

REMARKS

Claims 1-18 and 20-136 are pending in the application. All claims have been rejected. Amendments to the claims have been made to more clearly claim the patentable subject matter and are not related to reasons of patentability. Reconsideration is respectfully requested. All claims are now believed to be in condition for allowance.

Claim Rejections Under 35 U.S.C. § 103

Claims 1-6, 8-12, 18, 20, 22-24, 26-30, 32, 24-38, 44-46, 48-50, 52-54, 56, 58-60, 62-64, 66-70, 76, 77, 79-81, 83-86, 88-90, 92-96, 102, 103, 105-107, 109-113, 115-117, 119-123, 129-130, 132-134, and 136 are rejected under 35 U.S.C. § 103(a) as being unpatentable over U.S. Patent 5,926,813 to Chaudhuri et al. (Chaudhuri I) and U.S. Patent 6,223,171 to Chaudhuri et al. (Chaudhuri II). Claims 1, 27, 53, 54, and 110 are independent claims. Reconsideration is respectfully requested.

Claims 1 and 27

The Applicants claim a mechanism for finding an optimal set of indexes for a database, based on the database schema, workload information, and an index superset formed by a union of the current index set and a proposed index set. By starting with this limited superset of indexes, the claimed invention is able to focus the index verification process from early on. Statistics for the derived index set are collected from the database.

The proposed index set can be proposed, for example, by specialized tools, such as Oracle Expert, that use extensive rule bases to suggest particular indexes. Also, an experienced administrator can feed in a proposed index set. Starting from this limited superset (the union of the current index set and the proposed index set) saves a significant amount of resources and time to reach an optimal solution.

Chaudhuri I, on the other hand, starts with the set of all possible indexes, and then attempts to reduce the number of indexes and therefore index configurations for evaluation (Chaudhuri I, col. 6, ll. 62-67). Thus, Chaudhuri I's technique consumes more resources and time to reach an optimal solution than the Applicants' claimed invention. Furthermore, Chaudhuri I

does not analyze collected statistics, because the indexes of Chaudhuri I are "what if" indexes and, as such, have no associated collected statistics.

Chaudhuri I, column 7, lines 56-63, cited by the Examiner, discusses gathering statistics and storing index entries, but its statistics are artificially generated, not collected from the database, as recited in Applicants' Claim 1.

Applicants therefore respectfully request reconsideration and withdrawal of the rejection of Claim 1 as amended, in favor of allowance. Allowance of dependent claims 2-18 and 20-26, and dependent claims 55 and 56, follows from the allowance of Claim 1.

Claim 27 has been amended to recite the collected statistics and should now be allowable for the same reasons as Claim 1. Allowance of dependent claims 28-52 and new dependent claim 57 follows from the allowance of Claim 27.

Claims 53 and 54, and Claims 58-109

Like Claim 27, Claims 53 and 54 have been amended to recite the collected database statistics and should now be allowable for the same reasons as Claim 1.

Dependent Claims 58-109 correspond to Claims 2-18, 20-26, 55 and 56, and their allowance follows from the allowance of Claims 53 and 54, respectively.

Claims 6 and 32

An optimizer typically generates a cost, and an "execution plan," which is an optimal series of steps to be executed to fulfill execution of the subject workload. The Applicants' claimed invention, as recited in claims 6 and 32, evaluates the plan to generate at least one statistic.

For example, an execution plan for a given workload might indicate that table 1 is to be accessed using index a, that table 2 is to be accessed using index b and finally that x rows of information are to be returned. The indicated index utilization, for example, the fact that the plan uses index a, can contribute to the index usage statistics of the invention as recited in claim 6. The reason for this is that cost is not sufficient when determining tradeoffs as to whether or not to include an index in the solution. The Applicants have found that consideration of other factors

that can be gleaned from the execution plan, such as SQL statement frequency or index volatility, or even a user-specified SQL statement importance, can benefit the index selection process. These additional factors allow the mechanism of the claimed invention to make better choices for indexing solutions when, for example, the optimizer costs are similar.

Furthermore, with a high insert-intensive workload, the Applicants' invention can use the volatility information to guide the solution towards fewer indexes, possibly choosing a solution with a higher cost that minimizes the number of indexes.

Chaudhuri I, at col. 10, lines 22-25 and 53-60, cited by the Examiner, teach that an optimizer returns a cost and an execution plan. The cost is evaluated, but nothing in the passage, alone or in combination with the other references, teaches or suggests that the cost evaluation tool looks at and evaluates the plan itself.

The Applicants therefore respectfully request the reconsideration and withdrawal of the rejections of claims 6 and 32 in favor of allowance, regardless of the final disposition of their respective base claims. Allowance of their respective dependent claims, i.e., claims 7-19 and 33-45, would then follow.

Rejections of Claims 21, 47, 78, 104 and 131

Dependent claims 21, 47, 78, 104 and 131 have been rejected under 35 U.S.C. § 103(a) as being unpatentable over Chaudhuri I and Chaudhuri II, in view of U.S. Patent No. 6,021,405 to Celis et al. The Applicants respectfully traverse this rejection.

First, the Applicants disagree that it would be obvious to combine Chaudhuri I and II with Celis. Celis is directed to optimizing the execution of a query by transforming logical expressions into execution plans (Celis, Abstract). This is typically done automatically after submission of the query. (*See, e.g.*, Celis, col. 12, ll. 40-65.) Both Chaudhuri I and Chaudhuri II, on the other hand, are directed to reducing costs and time in selecting index sets (Chaudhuri I, Abstract and Chaudhuri II, Abstract). Indexes are simulated but not actually created. (Chaudhuri I, col. 6, ll. 46-50). Thus, as pertains to Chaudhuri I and II, queries are not executed, but rather are used in the analysis of various candidate index sets. The Applicants do not believe that one skilled in the art would be motivated to modify Celis's transformance of logical expressions with Chaudhuri's

simulated "what-if" indexes, or vice versa. The Applicants thus respectfully request the reconsideration and withdrawal of the rejections of claims 21, 47, 78, 104 and 131 under 35 U.S.C. § 103(a).

Even when combined, however, Chaudhuri I, Chaudhuri II, and Celis together do not teach the Applicants' claimed invention.

The Applicants' claimed invention uses a reduced workload of unique statements rather than a full workload during the evaluation process (Specification, p. 3, ll. 4-5, and Fig. 3A). The reduced workload compresses the set of SQL statements in the workload down to a smaller set that is relevant for index analysis. The Applicants have found that using a full, unreduced workload is not practical for real-world applications. Such analysis has been found to be very resource-intensive and time-consuming without the use of a reduced workload.

Celis, column 5, lines 37-44, on the other hand, discusses tracking rules to eliminate the generation of redundant expressions that can occur when rules (i.e., transformation rules which produce equivalent expressions – *see* Celis, col. 2, ll. 25-28) are applied multiple times to an expression. Celis says nothing about reducing a workload into unique statements, as recited in claims 21, 47, 78, 104, and 131. That is, Celis does not teach or suggest "the workload is reduced into unique statements," as recited in claims 21, 47, 78, 104, and 131. Therefore, these claims are allowable regardless of the final dispositions of their respective base claims.

Rejections of Claims 25, 51, 82, 108 and 135

Dependent claims 25, 51, 82, 108, and 135, each of which recites that the proposed index list is provided by an expert system, have been rejected under 35 U.S.C. § 103(a) as being unpatentable over Chaudhuri I and Chaudhuri II, in view of Lomet et al. The Applicants respectfully disagree. "Expert system" is a well-known term in the computer industry. An expert system is a computer program that simulates the judgement and behavior of a human or an organization that has expert knowledge and experience in a particular field. Typically, such a system contains a knowledge base containing accumulated experience and a set of rules for applying the knowledge base to each particular situation that is described to the program. Sophisticated expert systems can be enhanced with additions to the knowledge base or to the set

of rules (*See, e.g.*, http://whatis.techtarget.com/definition/0,,sid9_gci212087,00.html, last updated on July 30, 2001).

Although Lomet, as the Examiner suggests, discusses the use of expert systems to select an index, Lomet actually teaches away from a combination of an expert system with a system such as the Applicants' in which a proposed set of indexes is suggested for further evaluation. For example, in discussing prior work, Lomet states: "The class of tools that adopt an expert system like approach ... suffer from being disconnected from the query optimizer." (Lomet, page 83, left column). Lomet's tool, the "index tuning wizard," merely iterates through the space of hypothetical indexes. There is no teaching or suggestion that an expert system could be used to provide a proposed index set from which, along with a current index set, statistics can be generated, as recited in claims under rejection.

Chaurhuri I, column 5, line 57 to column 6, line 13, cited by the Examiner, discusses an index selection tool but does not suggest that this tool makes use of an expert system.

Therefore, claims 25, 51, 82, 108, and 135 are allowable regardless of the final disposition of their respective base claims. The Applicants respectfully request the reconsideration and withdrawal of the rejections of these claims in favor of allowance.

Rejection of Claims 7, 33, 65, 91, and 118

Dependent Claims 7, 33, 65, 91 and 118, all of which recite basing the execution plan on available access paths, have been rejected under 35 U.S.C. § 103(a) as being unpatentable over Chaudhuri I and Chaudhuri II, in view of U.S. Patent No. 5,924,088 to Jakobsson et al. The Applicants respectfully disagree.

The Applicants respectfully disagree that it would be obvious to combine Chaudhuri I and II with Jakobsson. Jakobsson is directed to using a cost/benefit analysis in selecting index access methods that include index access paths in response to a query containing a WHERE clause (Jakobsson et al., Abstract).

As discussed above, both Chaudhuri I and Chaudhuri II are directed to reducing costs in selecting index sets by simulating, but not actually creating the indexes. Thus, one skilled in the art would not be motivated to combine the index access methods of Jakobsson with the simulated

indexes because there would be no access paths for the nonexistent indexes of Chaudhuri I and II. The Applicants respectfully request reconsideration with respect to Claims 7, 33, 65, 91, and 118.

Rejection of Claim 31

Dependent claim 31, which recites "eliminating at least one index on a small table," has been rejected under 35 U.S.C. § 103(a) as being unpatentable over Chaudhuri I, Chaudhuri II, U.S. Patent No. 5,924,088 to Jakobsson et al. and Siegel et al. The Applicants respectfully disagree.

Jakobsson, column 12, line 66 to column 13, line 1, cited by the Examiner, teaches that the cost of searching a table without using an index is proportional to the size of the table. The purpose of the discussion appears to be to show how the cost improves when an index is added. However, Jakobsson does not teach or suggest, as the Examiner asserts, eliminating an index on a small table, for example, where "the table is sufficiently small that an index would never be a desirable access method." (Specification, p. 20, ll. 9-10).

None of the cited references, alone or in combination, teach or suggest "eliminating at least one index on a small table," as recited in claim 31. Therefore claim 31 is allowable. The Applicants therefore respectfully request the reconsideration and withdrawal of the rejection of claim 31, in favor of allowance.

Rejection of Claims 13-16, 39-41, 71-74, 97-100, and 124-127

Claims 13-16, 39-41, 71-74, 97-100, and 124-127, all of which recite deriving the cost of the execution plan from a resource use needed to execute the statement according to the execution plan, have been rejected under 35 U.S.C. § 103(a) as being unpatentable over Chaudhuri I and Chaudhuri II, in view of U.S. Patent No. 6,003,022 to Eberhard et al. Applicants respectfully request reconsideration.

The Applicants respectfully disagree that it would be obvious to combine Chaudhuri I and II with Eberhard. Eberhard merely discusses estimating execution costs based on simplified

transaction statements (Eberhard, Abstract). It does not address using these costs in connection with execution plans.

Thus, one skilled in the art would not be motivated to combine the cost estimation of Eberhard with proposed indexes of Chaudhuri I and II because the costs associated with indexes of Chaudhuri are index creation and maintenance costs, not statement execution costs as shown in Eberhard. (*See Chaudhuri I*, col. 12, ll. 20-35.) Therefore, the Applicants respectfully request reconsideration with respect to Claims 13-16, 390-41, 71-74, 97-100, and 124-127.

Rejection of Claims 17, 43, 75, 101, and 128

Claims 17, 43, 75, 101, and 128, all of which recite incorporating user-defined importance of the statement into the statistics, have been rejected under 35 U.S.C. § 103(a) as being unpatentable over Chaudhuri I and Chaudhuri II, in view of U.S. Patent 5,404,510 to Smith et al. Applicants respectfully request reconsideration.

The Applicants respectfully disagree that it would be obvious to combine Chaudhuri I and II with Smith. Smith addressees selecting existing or candidate indexes based on user-defined priority ranking. One skilled in the art would not be motivated to combine this index-selection mechanism of Smith with simulated “what-if” indexes of Chaudhuri I and II. Thus, the Applicants respectfully request reconsideration with respect to Claims 17, 43, 75, 101, and 128.

Rejection of Claims 55, 57, 61, 87, and 114

Claims 55, 57, 61, 87, and 114, all of which recite using index volatility as part of the collected statistics, have been rejected under 35 U.S.C. § 103(a) as being unpatentable over Chaudhuri I and Chaudhuri II, in view of Gurry et al. (Oracle Performance Tuning). Applicants respectfully request reconsideration.


As discussed, above, Chaudhuri I and II do not teach or suggest collecting and analyzing database statistics. It would not have been obvious to one skilled in the art to combine the volatility statistics of Gurry with the simulated indexes of Chaudhuri I and II because no volatility statistics exist for the “what-if” indexes of Chaudhuri I and II. Therefore, the Applicants respectfully request reconsideration with respect to Claims 55, 57, 61, 87, and 114.

CONCLUSION

In view of the above amendments and remarks, it is believed that all claims are in condition for allowance, and it is respectfully requested that the application be passed to issue. If the Examiner feels that a telephone conference would expedite prosecution of this case, the Examiner is invited to call the undersigned attorney.

Respectfully submitted,

HAMILTON, BROOK, SMITH & REYNOLDS, P.C.

By 
Rodney D. Johnson
Registration No. 36,558
Telephone: (978) 341-0036
Facsimile: (978) 341-0136

Concord, MA 01742-9133

Dated:

March 24, 2004